

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 709 779 A2

QUAD-55767
USSN 09/592,572

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
01.05.1996 Bulletin 1996/18(51) Int. Cl.⁶: G06F 11/14

(21) Application number: 95115752.8

(22) Date of filing: 06.10.1995

(84) Designated Contracting States:
DE FR GB

(30) Priority: 31.10.1994 US 332157

(71) Applicant: International Business Machines
Corporation
Armonk, N.Y. 10504 (US)(72) Inventors:
• Attanasio, Clement Richard
Peekskill, New York 10566 (US)
• Butrico, Maria Angela
Blauvelt, New York 10913 (US)

- Peterson, James Lyle
Austin, Texas 78759-5108 (US)
- Polyzois, Christos Alkiviadis
White Plains, New York 10606-1935 (US)
- Smith, Stephen Edwin
Mahopac, New York 10541 (US)

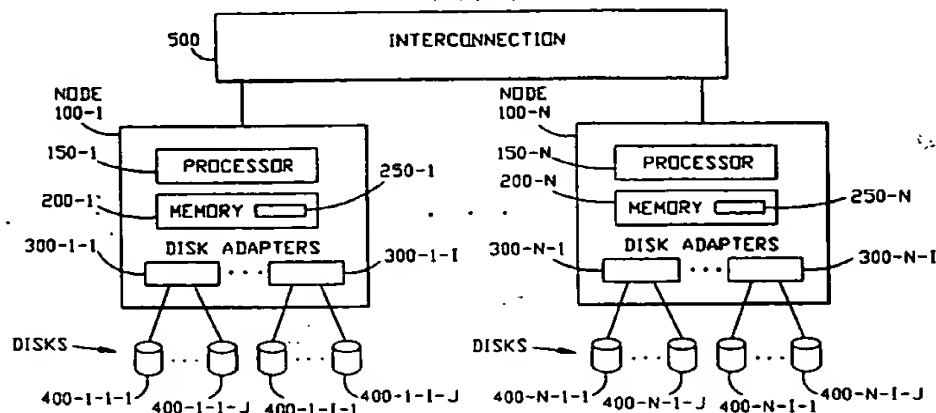
(74) Representative: Rach, Werner, Dr.
IBM Deutschland
Informationssysteme GmbH,
Patentwesen und Urheberrecht
D-70548 Stuttgart (DE)

(54) Virtual shared disks with application-transparent recovery

(57) A system and method for recovering from failures in the disk access path of a clustered computing system. Each node of the clustered computing system is provided with proxy software for handling physical disk access requests from applications executing on the node and for directing the disk access requests to an appropriate server to which the disk is physically attached. The

proxy software on each node maintains state information for all pending requests originating from that node. In response to detection of a failure along the disk access path, the proxy software on all of the nodes directs all further requests for disk access to a secondary node physically attached to the same disk.

FIG. 1



EP 0 709 779 A2

Description

I. Background of the Invention

a. Field of the Invention

This invention relates generally to a distributed computing environment. More particularly, it relates to a method for use in a cluster of processors, wherein each processor in the cluster can access any disk in the cluster.

b. Related Art

The availability of powerful microprocessors has made clusters an attractive alternative to monolithic systems. Applications that can partition their computation among several nodes can take advantage of this architecture, which typically offers better price-performance than the monolithic systems. Such applications include large scientific computations, database and transaction processing systems, decision support systems, and so on.

A microprocessor cluster consists of a number of separate computing systems, coupled with an interprocessor communication mechanism, such as a network or communications switch. Each computing system has its own processor, memory, and I/O subsystem, and runs a separate instance of the operating system. For maximum benefit, however, it is desirable for an application to be able to abstract from the specific computing system, and treat all nodes in a cluster as equivalent. This ability is sometimes called a "single system image."

A useful aspect of single system image is the requirement that the same I/O device resources be available to all processors in the cluster equally. This allows processing tasks to be freely moved between processors. Furthermore, it facilitates the development of parallel applications that adopt a data sharing model for their computation.

Many different approaches can be taken to providing the same I/O resources to all processors, preferably in a highly available fashion. Data replication is the simplest, especially for read-only data, but it increases cost (resources not shared) and presents difficulties when the information changes over time.

An alternative is to have devices that can be physically attached to many processors. For example, twin-tailed (dual ported) disks are common. It is possible to build four-tailed disks, and even eight-tailed disks, but they become increasingly expensive and difficult to operate.

In both of the above cases, each processor has independent access to the resources, so no action is necessary to provide continuous access to the data in case of processor and/or adapter failure.

Distributed file systems, such as NFS, AFS and DFS, abstract away from the specific I/O device to the services it is intended for and provide those services to

the processors in the cluster. This restricts the use of the device to those services, thus making it inappropriate for applications that are explicitly aware of the location of the data in the memory hierarchy. For example, a database system may rely on its own buffering and may want to arrange the data on disk in its own way, rather than rely on a file system to provide these services. In this case, direct access to the I/O device may be preferred by the application.

In terms of high availability, HA-NFS presents NFS clients with a highly available NFS server, but it relies heavily on the underlying network technology (IP address takeover) to provide critical functions that enable high availability.

II. SUMMARY OF THE INVENTION

It is an object of this invention to provide transparent recovery from node and/or adapter failure, in a system that allows processors in a cluster to share I/O devices without requiring that every I/O device be attached to every node, and allow applications running on surviving nodes to continue processing despite the failure. By transparent, it is meant that applications do not have to reissue any requests they had issued before the failure occurred.

Accordingly, the present invention provides a system and method for recovering from failures in the disk access path of a clustered computing system. Each node of the clustered computing system is provided with proxy software for handling physical disk access requests from applications executing on the node and for directing the disk access requests to an appropriate server to which the disk is physically attached. The proxy software on each node maintains state information for all pending requests originating from that node. In response to detection of a failure along the disk access path (e.g. in a node or disk adapter), the proxy software on all of the nodes directs all further requests for disk access to a secondary node physically attached to the same disk.

In a preferred embodiment the proxy software is embodied as a software layer that enables processors to access I/O devices physically attached to remote processors by defining virtual devices, intercepting I/O requests to those devices, and routing the requests (and data, for writes) to the appropriate server processor, to which the real device is physically attached. The server processor performs the actual I/O operation and returns a completion message (and data, for a read) to the originating processor. Upon receipt of the completion message, the originating processor notifies, accordingly, the process that had issued the request.

With twin-tailed disks, high availability can be achieved as follows. For a particular disk, one of the processors attached to the disk is designated as the primary server. During normal operation, I/O requests for the disk that originate anywhere in the cluster are sent to the primary server. If the primary server or its disk adapter fails, one of the other processors attached to the disk

becomes primary server for the disk, and the request routing information on each processor is changed, so that new requests are sent to the new primary server.

In the preferred embodiment, the server is totally stateless; the full state of pending remote requests is maintained on the client. Thus, in case of server or adapter failure, pending requests that had been issued prior to the failure can be re-issued by the client to the new server and the applications never see the failure.

III. BRIEF DESCRIPTION OF THE DRAWING

The present invention will be better understood by reference to the drawing, wherein:

FIG. 1 is an overall block diagram of a preferred embodiment of this invention;

FIG. 2 illustrates a preferred organization of twin-tailed disks;

FIG. 3 is a flow chart which shows the steps involved in processing a request at a client node;

FIG. 4 is a flow chart which shows the steps involved in processing a request at a server node;

FIG. 5 is a flow chart which shows the steps involved in recovery at the coordinator node;

FIG. 6 is a flow chart which shows the steps involved in recovery at the participant nodes; and,

FIG. 7 is a detailed block diagram of the memory resident logic and data related to the virtual shared disks.

IV. DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

FIG. 1 is a block diagram of a preferred embodiment of this invention, which incorporates the subsystem of recoverable virtual shared disks. It includes a collection (cluster) of independent computing nodes (henceforth nodes) 100-1 through 100-N.

Each node has a processor labeled 150-1 for node 100-1 and memory labeled 200-1 for node 100-1 (correspondingly 150-N and 200-N for node 100-N). Those skilled in the art will readily appreciate that each node could have separate memory or nodes could share memory.

The nodes can communicate via an interconnection 500. The interconnection 500 can be a switch, a local area network, shared memory, or any other kind of medium that allows the nodes in the cluster to exchange data.

Each node has a number of disk adapters labeled 300-1-1 through 300-1-I for node 100-1 (correspondingly, 300-N-1 through 300-N-I for node 100-N). Disks

labeled 400-1-1-1 through 400-1-1-J are attached to adapter 300-1-1 (correspondingly disks 400-N-1-1 through 400-N-1-J are attached to adapter 300-N-1). The number of disk adapters per node need not be the same for all nodes. Also, the number of disks per adapter need not be the same for all disk adapters. Some nodes may even have no disk adapters at all.

Disks which are shared by multiple nodes are addressed by a common name throughout the cluster, using the same programming interfaces used by a node to address a directly connected physical disk. This creates the illusion of the disks being physically connected to each node in the cluster. The software and programming interface which enables such accesses is referred to as a virtual shared disk.

Each processor's memory contains proxy logic and state data related to the virtual shared disks. The state data includes data of a type which is conventionally maintained by operating systems for physically attached disks (e.g., device state, device name, pending request information) as well as some additional information which will be described herein. This logic and related data is shown as block 250-1 for node 100-1 (correspondingly 250-N for node 100-N). Such a block for a node 100-K is shown in detail in FIG. 7. The proxy logic is shown as block 250-K-A in FIG. 7.

Disks which need to remain accessible in the event of a node or adapter failure, are attached to more than one adapter on different nodes. FIG. 2 illustrates the organization of a twin-tailed disk, where the disk 400-L-P-X is attached to adapter 300-L-P on node 100-L and adapter 300-M-Q on node 100-M.

During normal operation, for every disk, one of its tails is selected as the primary tail. Every node has a table (block 250-K-B in FIG. 7) that maps every virtual disk in the system to the node that holds the currently primary tail. The primary tail is the only tail used; the other tails of the disk are on stand-by.

Applications running on any node can issue I/O requests for any disk, as if all disks were attached locally. The logic for handling a request at the node of origin is shown in FIG. 3. When the request is issued (block 700), the aforementioned map, 250-K-B, is checked to determine which node has the primary tail (block 710). If the node of origin is also the server node (i.e., holds the primary tail), the request is serviced locally (block 715). If the server node is different from the node of origin, a request descriptor is sent to the server node (block 720). If the request is a write request (determined in block 730), the data to be written is also sent to the server (block 740).

In block 750 the request (whether read or write) waits for a response from the remote server. When the response arrives, if the request was a read (determined in block 760), the data that came on the network is given to the original request (block 770). If the request was not a read, the request completes (block 780).

The request descriptor includes the same type of data that an operating system would conventionally send

to a physical disk device driver (e.g. device name, offset, size of request, option flags, request type) as well as additional data such as the node of origin, a unique request identifier and the address of the primary node.

The logic for handling a request at the server node is shown in FIG. 4. A request (block 800) may have either been issued by a process running locally, or it may have arrived on the network from a remote node. In block 810, the access request is issued to the device. In block 820 the logic determines the source of the request. Upon I/O completion, if the request originated locally, the operation is complete and the originating process is notified. If the request originated at another node, in block 830 a response is sent back to the node of origin. If the request was a read, the data read is also sent. When the response arrives at the node of origin, the operation completes and the originating process is notified.

When a processor or adapter failure occurs, it is detected, in a conventional manner and all nodes are notified accordingly. Those skilled in the art will readily appreciate that various mechanisms (e.g., based on periodic health checks) can be used to detect failures.

In case of a node failure, the virtual shared disks affected are the ones for which the failed node was serving as the primary tail for the underlying physical device. In case of adapter failure, the virtual shared disks affected are the ones for which the primary tail of the underlying physical device was connected to the failed adapter. For all of the affected virtual shared disks, another tail of the underlying physical device is selected as the new primary tail. The selection can be made either statically (on the basis of a predetermined preference order) or dynamically (by a policy module that uses run-time information). The dynamic selection logic can be such that it attempts to achieve load balancing between the remaining active tails.

One of the nodes in the cluster is designated as coordinator. The coordinator is responsible for notifying all nodes in the cluster about the failure. The logic executed by the coordinator is shown in FIG. 5. The logic for each participant node is shown in FIG. 6. The coordinator is also a participant.

Upon failure detection (block 900), the coordinator broadcasts a message to all participants (block 910), telling them to suspend the affected virtual shared disks. Upon receipt of this message (block 1000), each participant suspends the affected virtual shared disks (block 1010). Suspension of a virtual shared disk means that the virtual device is marked as temporarily having no primary tail. Pending requests that had been sent to the failed server are saved by the client of origin in a queue provided for this purpose. Requests that arrive while a virtual device is suspended are also saved by the client of origin in the same queue.

After suspension of the affected devices, each participant sends an acknowledgement to the coordinator (block 1020) and waits (takes no further action with respect to the affected VSD) until it receives a resume message from the coordinator (block 1030). Other

processing is not affected. The coordinator waits for all participants to respond (block 920), and then broadcasts a message to all participants to resume the affected virtual shared disks (block 930). Upon receipt of this message, each participant resumes the affected virtual devices (block 1040). Resumption of a virtual device means that the node holding the selected new primary tail is recorded in the destination map for that virtual device. After resumption, the participant sends an acknowledgement to the coordinator (block 1050) and in block 1060 re-issues all pending requests (those pending prior to the suspension as well as those initiated during the suspension period) to the new server for the device. The coordinator collects the second-round acknowledgements from all nodes (block 940).

Agreement protocols other than the variant of two-phase commit we have described can be used to achieve the suspension and resumption of the virtual shared disks affected by the failure. Furthermore, in case of coordinator failure, another coordinator can be elected to perform the coordination of recovery.

A buffer with frequently or recently accessed data can be maintained in memory at the node that holds the primary tail for a disk. If the requested data is available in memory, the buffered memory copy is used; otherwise, a physical disk access must take place.

For twin-tailed or generally multi-tailed disks, any node physically attached to a disk can act as a server. Furthermore, any subset of the processors physically attached to the disk can act as servers simultaneously, i.e., there can be more than one primary tail active simultaneously. Nodes not attached to a particular disk can access that disk by shipping the request to any of its active servers. The choice of server can be made statically or dynamically.

It should be understood that the disk accesses being handled and rerouted in the event of failure are physical access commands rather than file system operations. In other words, in the present system, each node having a virtual shared disk issues commands (such as reads and writes to particular locations of the physical disk) to the disk device driver, as if the physical disk were directly connected to the node by way of a disk adapter. These commands are passed from the virtual shared disk software to disk driver software on a node directly connected to the primary tail (port) of the disk which, in turn, issues the command to the disk controller by way of the connected port.

Now that the invention has been described by way of the preferred embodiment, various modifications and improvements will occur to those of skill in the art. Thus, it should be understood that the preferred embodiment has been provided as an example and not as a limitation. The scope of the invention is defined by the appended claims.

Claims

1. A method for recovering from failures in a disk access path of a clustered computing system, comprising the steps of:

a) providing each given node of the clustered computing system with proxy logic for handling physical disk access requests from applications executing on the given node and for directing the disk access requests to a primary node to which the disk is physically attached; the proxy logic on each given node maintaining state information for all pending requests originating from that given node;

b) detecting a failure along the access path to a disk;

c) in response to detection of the failure, notifying the proxy software on all of the nodes to direct all further requests for access to the disk to a secondary node which is also physically attached to the the disk.

2. The method of Claim 1 wherein the access path includes the disk adapters and nodes to which the disk is physically attached and wherein the failure is detected in any of the nodes and the disk adapters.

3. The method of Claim 1 comprising the further steps of: in response to the detection of the failure, storing incoming access requests to the disk in a queue; and, rerouting the requests in the queue to the disk by way of the secondary node.

4. A clustered multi-processing system comprising a plurality of N nodes; a multiported disk having plurality of ports connected to M of the nodes, where N is greater than M; a failure detection mechanism, coupled to the nodes, for detecting failures along a disk access path between the disk and the nodes which are not physically connected to the disk; and, proxy logic on each of the nodes, coupled to the failure detection mechanism, for redirecting access requests to the multiported disk to another disk access path between the disk and the nodes not physically connected to the disk, when a failure is detected.

5. The system of Claim 4 comprising a queue, in each of the nodes, for storing incoming access requests to the disk; and, means for rerouting the requests in the queue to the disk by way of the another disk access path.

6. The system of Claim 4 wherein the failure detection mechanism includes means for detecting failures in any of the nodes physically attached to the disk and

in disk adapters, coupled to the disk, in each of the nodes physically attached to the disk.

7. A method for recovering from failures along a disk access path in a clustered computing system, comprising the steps of:

detecting a failure in the disk access path in the clustered computing system;

upon detection of the failure broadcasting a message to all nodes of the system having access to the disk;

in response to the message, suspending virtual shared disks on each node, saving pending requests that had been sent to the disk along the failed access path and saving requests that arrive while a virtual shared disk is suspended;

broadcasting a second message to the nodes to resume the affected virtual shared disks;

upon receipt of the second message, resuming the virtual shared disks at each node by recording the node holding a new primary tail in a destination map for that virtual device; and,

re-issuing all of the requests to the new primary tail.

FIG. 1




FIG. 2




FIG. 3

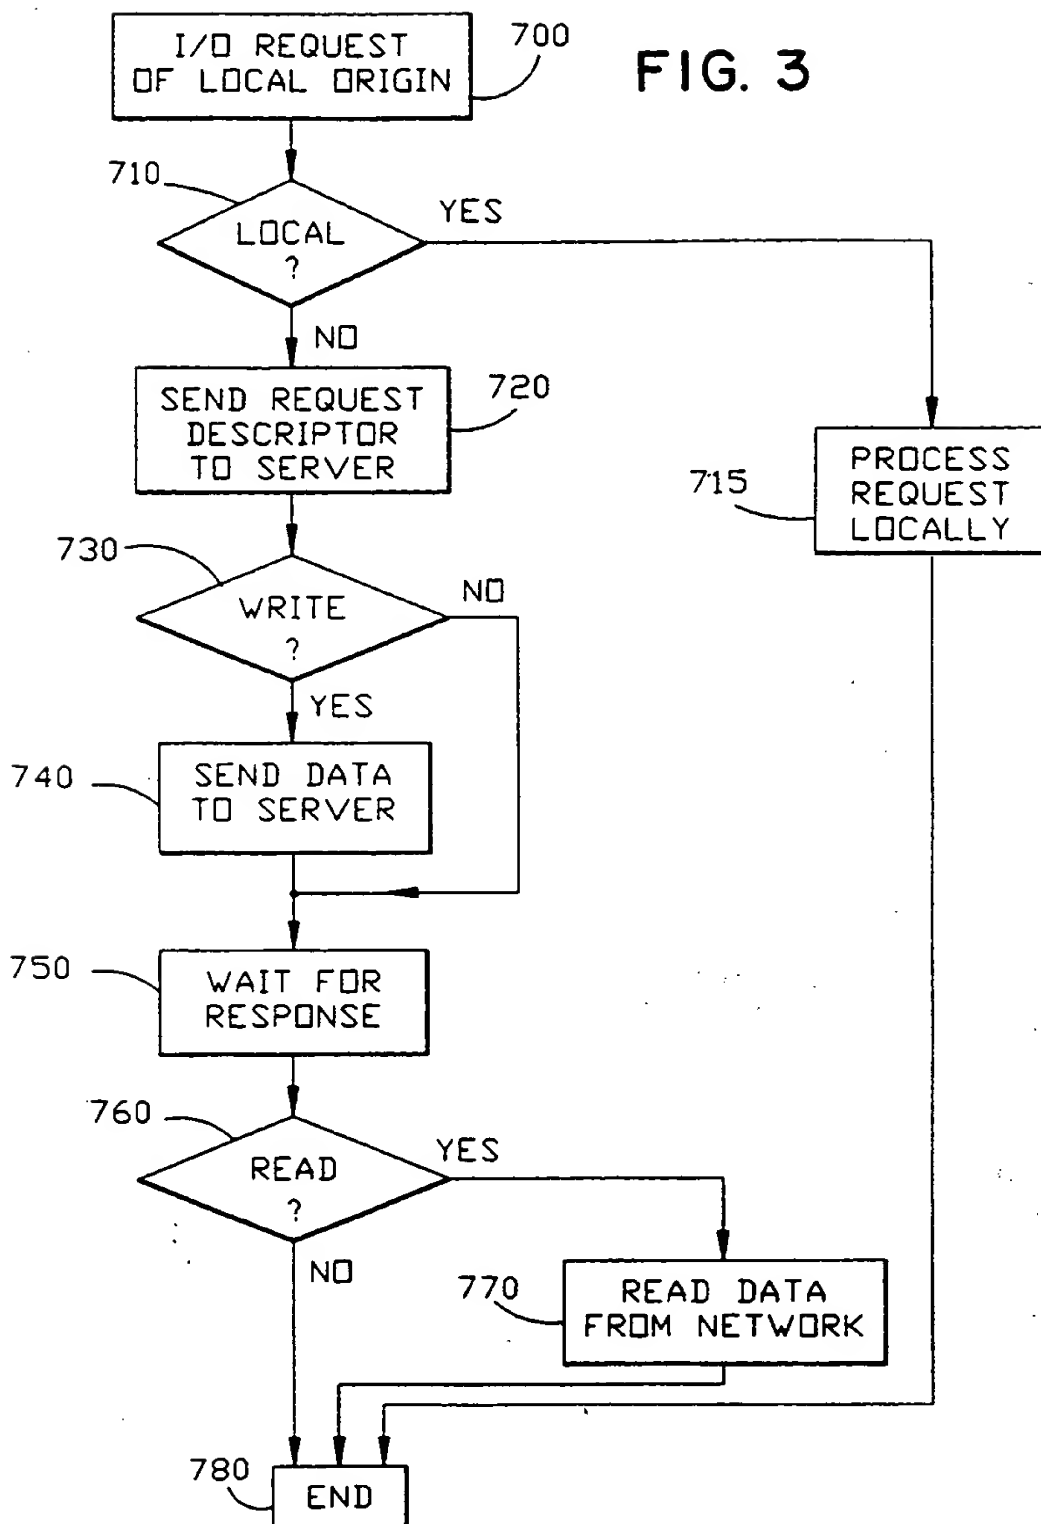


FIG. 4




FIG. 5




FIG. 6

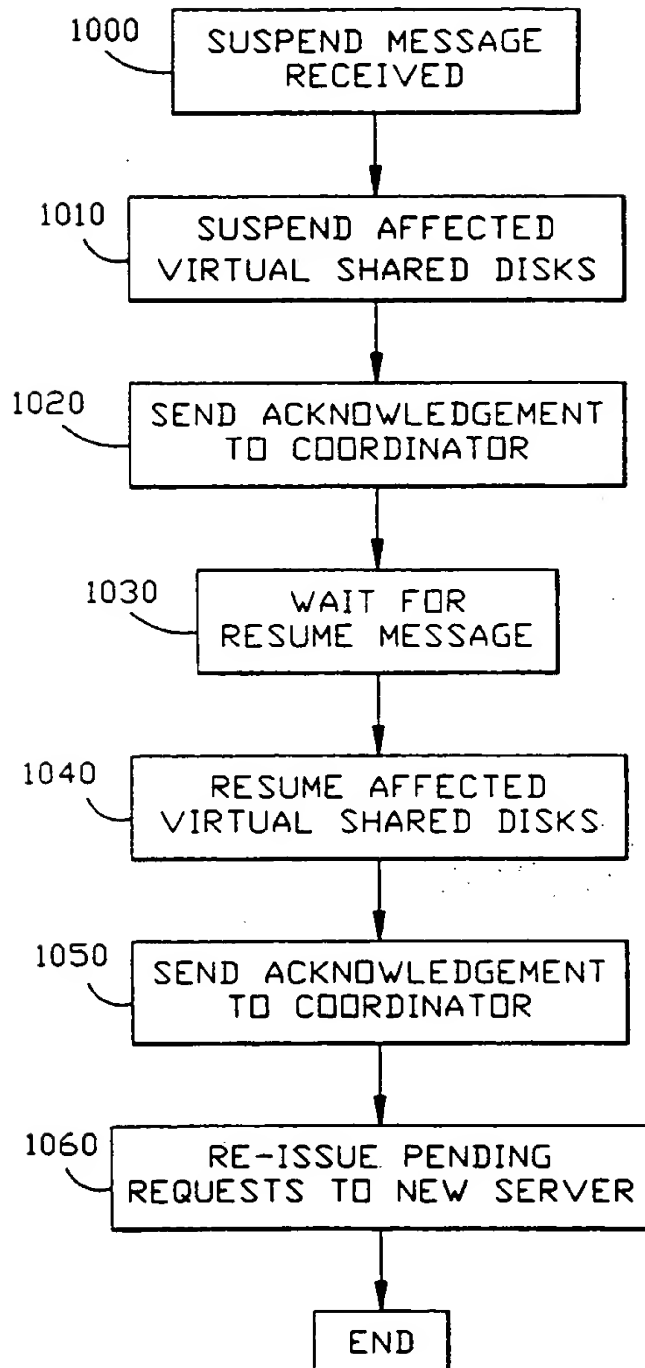




FIG. 7



THIS PAGE BLANK (USPTO)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 0 709 779 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
16.10.1996 Bulletin 1996/42

(51) Int. Cl.⁶: G06F 11/14, G06F 11/20

(43) Date of publication A2:
01.05.1996 Bulletin 1996/18

(21) Application number: 95115752.8

(22) Date of filing: 06.10.1995

(84) Designated Contracting States:
DE FR GB

(30) Priority: 31.10.1994 US 332157

(71) Applicant: International Business Machines
Corporation
Armonk, N.Y. 10504 (US)

(72) Inventors:
• Attanasio, Clement Richard
Peekskill, New York 10566 (US)
• Butrico, Maria Angela
Blauvelt, New York 10913 (US)

• Peterson, James Lyle
Austin, Texas 78759-5108 (US)
• Polyzois, Christos Alkiviadis
White Plains, New York 10606-1935 (US)
• Smith, Stephen Edwin
Mahopac, New York 10541 (US)

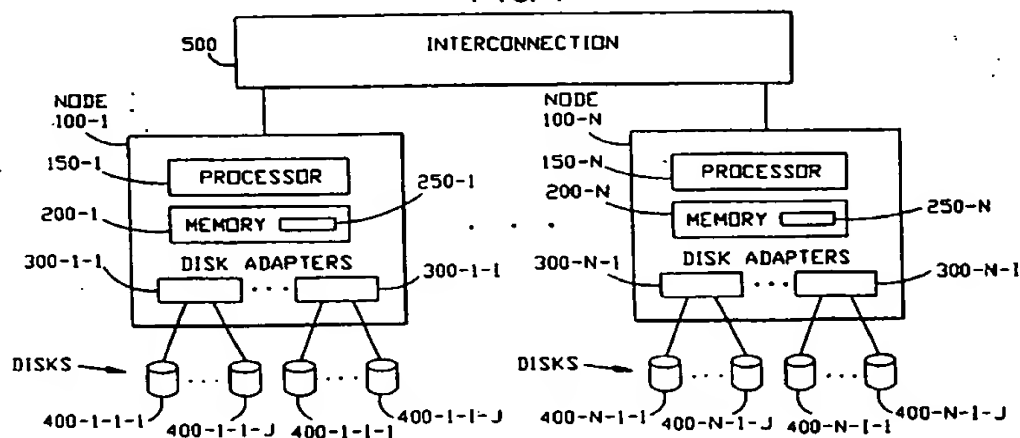
(74) Representative: Rach, Werner, Dr.
IBM Deutschland
Informationssysteme GmbH,
Patentwesen und Urheberrecht
70548 Stuttgart (DE)

(54) Virtual shared disks with application-transparent recovery

(57) A system and method for recovering from failures in the disk access path of a clustered computing system. Each node of the clustered computing system is provided with proxy software for handling physical disk access requests from applications executing on the node and for directing the disk access requests to an appropriate server to which the disk is physically

attached. The proxy software on each node maintains state information for all pending requests originating from that node. In response to detection of a failure along the disk access path, the proxy software on all of the nodes directs all further requests for disk access to a secondary node physically attached to the same disk.

FIG. 1



EP 0 709 779 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 11 5752

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 32, no. 2, July 1989, pages 378-380, XP000033461 "TAKEOVER SCHEME FOR CONTROL OF SHARED DISKS"	1-6	G06F11/14 G06F11/20
A	* the whole document *	7	
A	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 32, no. 11, 1 April 1990, page 168/169 XP000097659 "DYNAMIC STORAGE SUBSYSTEM PATH SWITCHING" * the whole document *	1,4	
A	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 36, no. 68, 1 June 1993, pages 375-377, XP000377422 "SHARED VIRTUAL DISK FOR A CLUSTER OF PROCESSORS WITH SEPARATE I/O DEVICES AND SHARED MEMORY" * the whole document *	1	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			G06F
The present search report has been drawn up for all claims			
Place of search BERLIN		Date of completion of the search 5 August 1996	Examiner Masche, C
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 (01.92) (P04.001)